

Estimating Attendance From Cellular Network Data

Marco Mamei

Dipartimento di Scienze e Metodi dell'Ingegneria
University of Modena and Reggio Emilia, Italy
marco.mamei@unimore.it

Massimo Colonna

Engineering & Tilab
Telecom Italia, Italy
massimo.colonna@telecomitalia.it

ABSTRACT

We present a methodology to estimate the number of attendees to events happening in the city from cellular network data. In this work we used anonymized Call Detail Records (CDRs) comprising data on where and when users access the cellular network. Our approach is based on two key ideas: (1) we identify the network cells associated to the event location. (2) We verify the attendance of each user, as a measure of whether (s)he generates CDRs during the event, but not during other times. We evaluate our approach to estimate the number of attendees to a number of events ranging from football matches in stadiums to concerts and festivals in open squares. Comparing our results with the best groundtruth data available, our estimates provide a median error of less than 15% of the actual number of attendees.

Categories and Subject Descriptors

G.3 [Probability and statistics]: Time series analysis; H.3.3 [Information Search and Retrieval]: Retrieval Models; I.5.2 [Design Methodology]: Pattern Analysis

General Terms

CDR, attendance estimation, mobility patterns

1. INTRODUCTION

The widespread diffusion of mobile phones and cell networks provides a practical way to collect geo-located information from a large user population. The analysis of such collected data is a fundamental asset in the development of pervasive and mobile computing applications, including location-based services, traffic management, urban planning, and disaster response [1, 2, 3, 4, 5, 6, 7].

In this work, we explore the use of anonymized Call Detail Records (CDRs) from a cellular network to estimate the number of attendees to large events happening in the city.

Each CDR contains information such as the time a mobile phone accesses the network (e.g., to send/receive calls and text messages), as well as the identity of the

cell tower with which the phone was associated at that time. CDRs can serve as sporadic samples of the approximate locations of the phone's owner.

On the basis of such location samples, we try to understand if a user was attending a given event and estimate the number of attendees on that basis.

While in some contexts, the number of participants can be deducted also by other means (e.g., ticketing information), there are many scenarios in which counting the attendance is problematic (e.g., events held in open squares, parades, flash-mobs) and an estimate on the basis of cellular network data is highly valuable.

Estimating events' attendance has a number of practical and useful applications.

On the one hand, it is an important information for the local government and organizers in that it is at the basis of event's planning and resource prioritization. In addition, since CDRs allow to track the movements of individual users, it is possible to understand where attendees come from and where they go after the event. This naturally supports traffic and road management.

On the other hand, such kind of information, can support advertisement systems [8] by providing accurate audience measurements. Also in this case, the possibility of tracking users would open to advanced applications for the provisioning of highly personalized advertising and marketing schemas. Despite users' hashed ids do not allow to identify the real person behind a phone, this opens a number of privacy concerns. While some research addressing such concerns exist [9, 10], we will not tackle privacy problems in this paper focusing on the attendance estimation problem only.

While a number of existing works deal with the problem of discovering and analyzing events on the basis of cellular network data (see Related Work section), the problem of actually estimating the number of attendees is largely unexplored. In particular, to the best of our knowledge, there are not published results of the accuracy of attendance estimation using CDRs.

The goal of this paper is to present such an estimation procedure. In particular, in Section 2 we present a naive approach to estimate the attendance and illus-

trate why it does not work properly. In Section 3 we present our methodology. In Section 4, we evaluate our approach to estimate the number of attendees to football matches in stadiums, in which reliable groundtruth data were available. Section 5 discusses how to improve performance on the basis of the knowledge of multiple events in the area. Section 6 presents related work. Eventually, Section 7 concludes and discuss some future avenues for improvement.

2. NAIVE APPROACH

Before illustrating the proposed methodology, we want to show the main problem that complicates the task of estimating the number of attendees.

A naive approach to address such an issue would be to just count the number of users who generate CDRs in cells covering the event’s location area during the event time. In particular, we tried to apply the naive approach to estimate the number of attendees to football matches in two different stadiums in Turin, Italy. We defined the area associated to each stadium as a circle centered in the stadium with a fixed radius of 100m. Then, we record all the CDRs produced in the network cells that overlap with the stadiums’ area at the event time. We then counted the number of individual users.

Figure 1 illustrates the result. The graphs represent the hourly count of users in the area associated with the stadiums (Stadio Olimpico on the left, Juventus Stadium on the right). We also highlighted football matches taking place in the stadiums with also groundtruth estimates for the number of attendees.

It is rather easy to see that the naive approach is highly ineffective. For example, the match that happened on March 12, 2012 at the Stadio Olimpico is reported to have 21453 attendees and a CDR users’ count with a peak of about 3700. In contrast, the match that happened on March 20, 2012 at the Juventus Stadium is reported to have a double number of attendees (40045), while a CDR users’ peak of about one-sixth (600).

The problem with these numbers is not in the discrepancy between groundtruth and CDR counts. This can be naturally explained by the fact that not all the users use the phone during the match, and by the fact that not all of them adopts the same carrier providing the data for this analysis.

The problem is in the negative correlation between groundtruth and CDR counts: large events (happening at the Juventus Stadium) appear to be smaller than “small” ones (happening at the Stadio Olimpico).

The reason for such a negative correlation can be easily found in the geography of the city. Stadio Olimpico is right in the city center. Juventus Stadium is in the suburbs. Accordingly, while network cells around Juventus Stadium are likely to measure CDRs coming from the stadium itself, network cells around Stadio




Figure 1: Hourly count of users generating CDRs in the area associated with the stadiums (Stadio Olimpico on the left, Juventus Stadium on the right). The problem is in the negative correlation between groundtruth and CDR counts: large events appear to be smaller than “small” ones.




Figure 2: Correlation result using the naive approach. It is easy to see that there is almost no correlation ($r^2 = 0.016$) among CDR count and groundtruth.

Olimpico overlap with a number of other relevant places and businesses in the city center thus inflating the result.

More in general, Figure 2 shows correlation results – using the naive approach – for a number of events covered by our dataset. Each point represents an event: the x-coordinate is the CDR estimate for attendance, while the y-coordinate is the groundtruth attendance. It is easy to see that there is almost no correlation ($r^2 = 0.016$) between the two estimates, so the naive approach is highly ineffective. Our goal is to identify a mechanism to create a strong *positive* correlation between groundtruth and CDR counts. Once this result is achieved, a simple linear regression can scale up CDR counts to the actual attendees estimate.

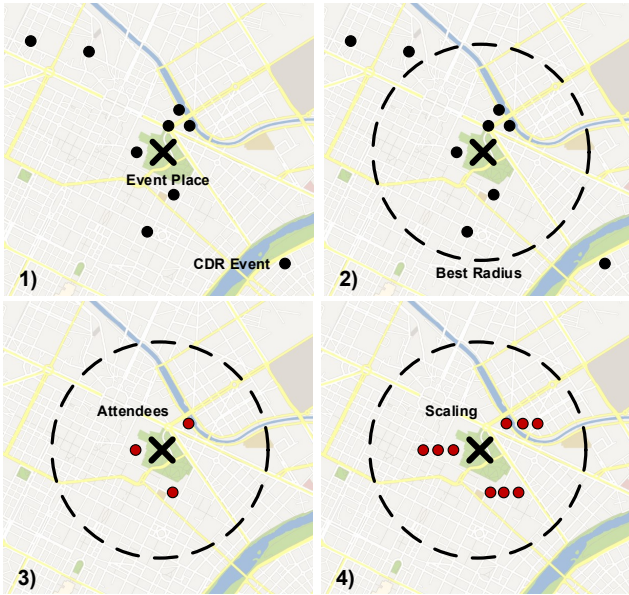


Figure 3: Proposed methodology to estimate event’s attendance. 1) We collect CDRs generated around the event area is selected. 2) We compute the radius best describing the event area. 3) The number of users who generate CDRs at the event time, but who do not (usually) generate CDRs at other times is recorded. 4) This number is then scaled according to a linear regression to find the actual attendance estimate.

3. METHODOLOGY

To overcome the above limitations, we developed a specific methodology to deal with attendance estimation (see Figure 3). In particular: (1) We collect all the CDRs generated around the event area. (2) We identify the radius within which are all the cells whose traffic can be associated to the area where the event takes place. (3) On the basis of the identified cells, we count the number of users who generate CDRs at the event time, but who do not (usually) generate CDRs at other times. Finally, on the basis of such data from a number of events, we set up a linear regression to estimate the number of attendees.

In the following subsections we describe in detail the above steps.

3.1 CDR Data

We obtained a large set of mobility data from an Italian telecom operator. In particular, we analysed data from two regions of Italy (Piemonte and Lombardia inhabited by about 15 millions people), spanning 16 months (March 2012 – June 2013) during which we analysed several events ranging from football matches

User	MMC	Time	Cell	Coord	Radius
3dd2b	222	7346286	123	(41.2,13.9)	450

Figure 4: Structure of our CDR dataset. Every time a user send or receive calls and text messages we generate one CDR with information about the user (hashed) id, the MMC (Mobile Country Code), the timestamp of the CDR, the code of the cell tower and the coordinates and coverage radius of the cell tower.

in stadiums to concerts and festivals in open squares.

Mobility data is obtained from Call Detail Records (CDRs) and Mobility Management (MM) procedure messages (i.e., IMSI attach/detach and Location Update) [11]. CDRs are routinely collected by cellular network providers for billing purposes. A CDR is generated every time a phone places or receives voice call or a text message. The IMSI attach/detach procedure marks the phone as attached/detached to the network on power up/power down of the phone or SIM inserted/removed. Location updates are messages exchanged for keeping the network informed of where the phone is roaming. CDR and MM messages are read on network interfaces through specific probes and also contain the identity of the phone, the identity of the cell through which the phone is communicating and the related timestamp. As MM messages, for the purposes of our study, contain the same information as CDRs, for simplicity of writing we will refer to all these data as CDRs.

In the context of this work, all this information serves as sporadic samples of the approximate locations of the phone’s owner. Specifically, the user’s location is given in terms of the cell network antenna the user was connected with. The area covered by a given antenna sector can be approximated by a circle with a given center and radius. In Figure 4 it is shown the structure of a CDR. Each record comprises a user (hashed) id, the MCC (Mobile Country Code) representing the country where the SIM card has been registered, the timestamp of the CDR, the code of the cell tower and the coordinates and coverage radius of the cell tower. Thus, the spatial resolution of CDR localization is the cell radius. Similarly to [12], in our work we take into considerations different sectors for different antennas. Each sector is referred to as an individual cell and approximated with a circle.

It is worth noticing that differently from a number of other works we do not estimate the coverage of a cell network by using Voronoi tessellation. We stick to the simpler representation of a cell being represented by a circle with a given center and radius. In [13], it is shown that the approach do not change the user’s location accuracy.

Figure 5 illustrates some key statistics of our data.




Figure 5: (left) Daily average number of CDRs produced for a given percentile of users. (right) Radius of gyration for a given percentile of users.

Figure 5-left illustrates the daily average number of CDRs produced for a given percentile of users. While the average number of CDRs per day is rather limited, we monitor a large user population comprising more than 4 million persons. In addition, as discussed in Section 3.3, CDRs are not evenly spread across all the days and across the 24 hours. So, we actually have more location samples in the time frame where events actually happen.

Figure 5-right illustrates the radius of gyration for a given percentile of users. The radius of gyration is a synthetic parameter describing the spatial extent of user traces. It is defined as the deviation of user positions from the corresponding centroid. It is given by: $r_g = \sqrt{\frac{1}{n} \sum_{i=1}^n (p_i - p_{centroid})^2}$ where p_i represents the i^{th} position recorded for the user and $p_{centroid}$ is the center of mass of the user’s recorded displacements obtained by: $p_{centroid} = \frac{1}{n} \sum_{i=1}^n (p_i)$. It is possible to see that almost half of the user are urban dweller with r_g less than 10Km. Users in the (50th-75th) percentiles can be associated to urban commuters as the diameter of peri-urban areas of main cities in the region is about 25-30Km. Users beyond the 75th percentile are associated to commuters travelling region-wide.

3.2 Best Radius

As discussed in Section 2, determining the cells that are relevant for the events generated in a given area is a fundamental task. Otherwise it is possible that the cells being considered will include CDRs actually produced elsewhere, or will miss CDRs that were actually produced in the proper area.

To tackle this problem, we model the event area as a circle with center c - where the event takes place, and with radius r . A cell with center b and radius rc is considered relevant for the event if: $dist(c, b) < r + rc$. Where $dist$ is the geographic distance between the points. In other words, a cell is relevant if it overlaps with the circle representing the event area.




Figure 6: Identification of the best radius to model the event area. If the radius is too large (top) the events’ structure cannot be identified properly. With a proper value of the radius (bottom) outlier in the CDR counts correspond to the events.

The problem of determining the relevant cells is thus shifted to the problem of identifying a proper radius r for the event area. It is important to notice that we could also select $r < 0$ to impose the fact that a cell has to overlap to the center of the event by a certain amount to be considered as relevant.

To solve this issue, our approach starts from the basic consideration that the plot of the number of CDRs generated from the event area should have a spike (i.e., an outlier) when the event takes place, as the events – we are interested in – will typically attract a large number of people.

For example, Figure 6 illustrates the z -score for the hourly count of users producing CDRs around a stadium (Stadio Silvio Piola, Novara, Italy). In the top graph, the stadium area is modeled as a circle with radius $r = 500m$. In the bottom graph, the stadium area is modeled as a circle with radius $r = -300m$ (see above discussion on negative radii). It is easy to see that adopting $r = 500m$ fails to capture the events’ structure in that events are not clear outliers. On the contrary with $r = -300m$ it is possible to precisely identify events (i.e., all the events have values larger than 3).

In this context, $r = -300m$ would be a suitable radius to describe the event area. This is probably due to the fact that the stadium is close to other relevant places and businesses. Taking large values of r bias the CDR count by considering also CDRs generated in these other places. Instead, a low value of r selects

only relevant CDRs. It is also possible to see that the outlier associated to the event on 29/4/2012 is readily visible even with $r = 500m$. The football match that happened on that date, in fact, attracted almost the double of people (17650 persons vs. stadium's average of 9370). Such an event would be better represented by a larger radius (the more the people, the more the cells nearby the stadium get saturated and rely the network connection to farther cells).

On the basis of the above considerations, we developed an approach to identify the best radius describing the event area. For each event happening at a location with center ec starting at time st and ending at et , we propose the the following approach. For the sake of clarity, we present the approach in two different steps.

STEP 1.

1. For different values of r in r_{min}, r_{max} , we extract the CDRs in the event area ($cdr[]$).
2. For each r_k , we compute the hourly count of users who generate CDRs in the area during the event time. We call x_k such a count.
3. We then compute the z -score of the x_k values in the event time frame. More in detail, we computed the hourly count of users who generate CDRs in the area during the event time, but in i days before the event (we considered 6 days before). We then computed the mean μ_k and standard deviation σ_k of this count. On this basis we computed the z -score $z_k = (x_k - \mu_k)/\sigma_k$. The result is a series of values z_k measuring how extreme the CDR count were during the event (considering given radius r_k).

```

Data:  $cdr[], ec, st, et$ 
Result:  $z[]$ 
forall the  $r_k \in [r_{min}, r_{max}]$  do
   $x_k = countUsers(cdr[], ec, r_k, st, et)$ 
  forall the  $i \in [0, 6]$  do
     $y_{ik} = countUsers(cdr[], ec, r_k, st - i \cdot days, et - i \cdot days)$ 
  end
   $\mu_k = mean_i(y_{ik})$ 
   $\sigma_k = sdt.dev_i(y_{ik})$ 
   $z_k = (x_k - \mu_k)/\sigma_k$ 
end

```

Algorithm 1: Radius Extraction - Step 1

Algorithm 1 presents a more formal description of the approach.

The result is a graph showing for each r_k how much the area had an unusually high number of people during the event. Figure 7 shows the result for two events. It is possible to see that once the area is properly identified,




Figure 7: Graph showing for each r_k how much the area had an unusually high number of people during the event.

the z -score clearly identifies that something unusual is taking place there ($z_k = 3.7$ with a radius of about 300m for the event on the left, $z_k = 2.2$ with a radius of about -200m for the event on the right).

STEP 2. On the basis of the graph showing the average z -score for different radii, we have to identify the actual *best* radius. Contrarily to a naive approach, selecting the radius associated to the maximum z is not an effective option. This approach would be strongly biased to small radii that always exhibit large z -scores. In fact, even if the event area is large, any (smaller) area contained in there would have a z -score that is likely to be higher than the whole event area, as it comprises only those cells that are really in the middle of the event. Accordingly, we adopted the following solution. See also Algorithm 2.

1. For each r_k , we normalize the z_k values by a factor representing the event area. The idea is that a large z over a small area around the event's location should be favoured with respect to a large area possibly comprising also other events. In particular, we divide each z_k by the sum of the radii of the network cells associated with the r_k area. More formally, calling nc_k the set of the network cells within the event area defined by r_k , and calling $nc.r_i$ their radii, the our normalized z value \hat{z} is computed as: $\hat{z}_k = z_k / \sum_{i \in NC_k} R_i$
2. Finally, we compute the best radius as the average of the r_k values weighted by the associated normalized z -scores.

```

Data:  $z[]$ 
Result:  $bestR$ 
forall the  $r_k \in [r_{min}, r_{max}]$  do
   $\hat{z}_k = z_k / \sum_{i \in nc_k} nc.r_i$ 
   $bestR = \frac{\sum_k r_k \cdot \hat{z}_k}{\sum_k \hat{z}_k}$ 
end

```

Algorithm 2: Radius Extraction - Step 2

3.3 Attendance Estimator

Once the event area has been identified, we need a mechanism to count precisely the number of users who attended the event. Since we do not know what the user was doing in the event area, we estimate the probability of the user presence as proportional to the fraction of time in which the user was there during the event, and inversely proportional to the fraction of time in which the user was there outside of the event time [14]. This latter point is important to eliminate users that live or work in the event area and so are in there independently of the event.

As a first step, we tried to characterize the individual calling activity and verified that it is frequent enough to allow monitoring the users' location with a fine enough resolution. For each user, we measured the inter-CDR time - i.e., the time interval between two consecutive network connections (similar to what has been done in [15, 1]). Focusing on a given event (e.g., a football game held at the Juventus Stadium in Turin on March 20 2012), we performed some measures. The average inter-CDR time measured for the population of possible attendees (users who generate at least one CDR in the event area during the event time) was 241 minutes. This number is large because it considers the whole daily lives of that users, thus also spanning night gaps. We also measured the average inter-CDR time considering only CDRs generated during at the event time. With that assumption the average inter-CDR time reduces to 52 minutes.

Because the distribution of inter-CDR times for a user spans several temporal scales, we further characterized each calling activity distribution during the event time by its first and third quartile and the median. Figure 8 shows the distribution of the first and third quartile and the median for all the possible attendees. The arithmetic average of the medians is 64 minutes (the geometric average of the medians is 51 minutes) with results small enough to detect changes of location where the user stops for about 2 hours.

Such a time frame should be compatible with the duration of a lot of the events of interest. We also verify that the above figures are consistent considering also other events.

On the basis of this analysis, we developed the following approach. We extract CDRs of all the possible attendees to an event, i.e., all the users that generate at least a CDR in the event area at the event time. Then, for each user:

1. We compute the user's average inter-CDR time iet in the daily hours in which the event takes place. We also compute the time of the *first* and of the *last* CDRs produced in the event area during the event time.




Figure 8: Characterization of individual calling activity for the population possibly attending an event, in terms of time between two network connections at event time. Graphs show the distributions of the median (red), first quartile (blue), and third quartile (green) of individual inter-CDR time

2. We compute the fraction of time in which the user is at the event. as:

$$f1 = \frac{|last-first+iet|}{eventduration}$$

3. We then compute, in the same way as before, the fraction of time in which the user is in the event area in a period spanning d days before the event (in our experiments we usually set $d = 6$). In particular, we compute the time of the *first* and of the *last* CDRs produced in the event area in the d days.

$$f2 = \frac{|last-first+iet|}{d \cdot days}$$

This represent the fraction of time in which the user is in the event area without the event.

4. We compute the probability of the user being at the event as

$$p = f1 \cdot (1 - f2)$$

For example, if the user was at the event for the whole event duration and (s)he never visited that area otherwise, then $p = 1$. Viceversa, if the user is always in the event area $p = 0$ because (s)he is likely to be there for other reasons than the event.

We then add all such probabilities p together to obtain a raw attendance estimator of the event. It is worth noticing that, in contrast with other approaches, we do not set a threshold to decide if a user was present or not. By adding the users' probabilities, it might happen that 2 users who attend the event with 50% probability are considered as 1 user attending with 100% probability. See Algorithm 3.

3.4 Linear Regression

The above estimator is typically much lower than the actual attendance. This can be naturally explained by

```

Data:  $cdr[], bestR, ec, st, et, d = 6$ 
Result:  $attendance$ 
 $candidates[] = usersIn(ec, bestR, st, et)$ 
forall the  $c_i \in candidates$  do
     $iet = avg - inter - CDR - time(c_i, cdr)$ 
     $first = timeFirstCDR(ec, bestR, st, et)$ 
     $last = timeLastCDR(ec, bestR, st, et)$ 
     $f1 = \frac{|last - first + iet|}{eventduration}$ 
     $first = timeFirstCDR(ec, bestR, st - d_{days}, et)$ 
     $last = timeLastCDR(ec, bestR, st - d_{days}, et)$ 
     $f2 = \frac{|last - first + iet|}{d_{days}}$ 
     $p_i = f1 \cdot (1 - f2)$ 
end
 $attendance = \sum_i p_i$ 

```

Algorithm 3: Attendance estimator

the fact that not all the users will use the phone during the event, and by the fact that not all of them adopts the same carrier providing the data for this analysis. In any case, as we will show in the next section, it has a strong *positive* correlation with groundtruth headcounts. Accordingly, a simple linear regression can scale up the above count to the actual attendees estimate.

Rather than more complex regression algorithms, we applied linear regression for two main reasons:

1. The goal of this work is to show that events' attendance can be *measured* by CDRs coming from the cellular network. If this is true, then an estimator based on CDR needs only to be scaled up to provide good results. More complex regression algorithms could hide shortcomings of the CDR estimator that we want instead to analyze.
2. The number of events for which we have groundtruth information is limited. Accordingly there is a notable risk of overfitting. Regression mechanisms more complex than linear regression would be even more susceptible to this problem.

More in detail, we assume the availability of a *training* set of events to be used to fit the parameters of the linear regression. The resulting coefficients are then used to scale CDR estimates of attendance in a *testing* set of events. The combination of all the above steps produces the final estimate of the number of attendees. In the next section, we conduce some experiments to assess the performance of our approach.

4. ANALYSIS AND RESULTS

As already introduced, to test the performance of the presented methodology we try to estimate the number of attendees to several events ranging from football matches in stadiums to concerts and festivals in open squares. The analysis spans large events with ground




Figure 9: Diagram showing best radius results for different places. It is possible to see that different events in the same place can be represented by different radii

truth attendance of more than 80000 persons to smaller ones with a ground truth attendance of less than 2000 persons. Overall, we considered a dataset comprising 43 events.

To take into account the fact that a number of CDRs might happen before and after the event, we set the event starting-time two hours before the official kick-off, and the event end-time two hours after the end of the event.

4.1 Best Radius

In this first set of experiments we report the radius that best capture the dynamics of a given event. We run the algorithm described in Section 3.2. Specifically, we varied r in $r_k \in [-500m, 1500m]$ with a $100m$ step. The result is a set of $NR = 21$ radii to be tested. Figure 9 illustrates the obtained results for different event areas under analysis (on the x-axis we indicate an id associated to different event areas – e.g., 1 = “a stadium in Bergamo, Italy”) It is possible to see that the same event area may be best represented with different radii depending on the specific event considered. This is rather natural, as the more people attend the event, the more the cells nearby the stadium get saturated and rely the network connection to farther cells. Accordingly, larger events (even in the same location) tend to be associated to larger radii.

4.2 Attendance Estimate

In this set of experiments we actually estimate attendance for multiple events. First we use the algorithm described in Section 3.3 to obtain a CDR count proportional to the attendance estimate. Then we scale that number with a linear regression. Specifically, for each event to be analyzed, we considered as a training set all the events happening in stadiums (leaving

out the considered event, if present). We use the estimated attendance of such events and the associated groundtruth attendance information to fit the parameters of a linear regression. We use events in stadiums as training set as they are typically associated with better groundtruth estimates (derived from ticketing information). We then scale the CDR count with the linear regression to obtain the final estimate. Specifically, we report results using different kinds of linear regression:

1. **Standard linear regression.** In this approach, we consider the whole training set, create a linear regression model fitted by minimizing sum of squared errors, and use the model parameters to scale predicted attendance count.
2. **Piecewise linear regression.** In this approach, for each testing sample, we consider the n closest samples in the training set, create a linear regression on that n points, and use it to scale that predicted testing sample. In our experiments we empirically set $n = 6$.
3. **Range linear regression.** We also conducted some experiments separating the events with an attendance below and above 10000 persons. This can be interpreted as a trade-off between global and piecewise regressions: we fit one regression for small (< 10000 persons) events, and another for large events (≥ 10000 persons).

Figure 10(top-row) illustrates the result of the different regressions between groundtruth data and our attendance estimator. Other than visually, we verified that in the case of linear regression (left plot), the results exhibit a Pearson correlation $r = 0.87$ and a coefficient of determination $r^2 = 0.76$ indicating a strong positive correlation between the results. In the case of piecewise regression (center plot) summarizing a single correlation coefficient is problematic. However, it is possible to see a good fit of the data. In the case of range linear regression (right plot), $r = 0.65/r^2 = 0.42$ for small events (< 10000 persons), $r = 0.93/r^2 = 0.86$ for large events (≥ 10000 persons), indicating offering weak results for small events, while strong correlation for large ones.

In all the plots, confidence intervals for the regression is depicted with a gray area.

Figure 10(bottom-row) illustrates mean/median absolute error between estimated attendees and groundtruth, and mean/median percentage error (absolute error divided by groundtruth). The gap in errors between mean and median indicates that the distribution of error is skewed (in the case of linear regression, skewness = 3.10, in the case of piecewise linear regression skewness = 2.69, in the case of range regression, skewness = -0.6/3.3 for small and large events respectively). This

is due to the fact that even small errors in the order of 1000 person would be very high in events with 2000 attendees (50% error) thus notably increasing the mean error.

To better quantify this behavior, Figure 11 shows error distribution with regard to groundtruth attendance (top-row) and the error CDF (bottom-row). The graph shows results for linear regression (left), piecewise linear regression (center), range regression (right). Looking at the graph, it is easy to see the skewness effect described above. In all the regressions, rather expectedly, the approach presents large errors for small events, while small errors for large events.

In summary, it is possible to see that the use of the described approach produces rather good estimates of the number of attendees. It is easy to see that results are better in large events where a limited absolute error has a small impact in the overall percent error. In general, we found that the proposed approach starts producing consistent good results for events larger than 10000 attending persons. Considering only those events with an attendance greater than 10000, Pearson correlation jumps to 0.93. Linear regression’s mean %error drops to 22% and median %error drops to 15%. Similarly, piecewise linear regression’s mean %error drops to 16% and median %error drops to 13%.

4.3 Unstructured Events

The dataset of events used for the experiments comprises two kinds of events: (i) “structured” events, like concerts and football matches, for which some sort of entrance policy (e.g., entrance gates) allow to obtain reliable estimate on the number of attending persons. (ii) “unstructured” events happening in open squares or parks for which no entrance policy is enforced. The analysis of this latter kind of events is problematic because it is very difficult to obtain reliable groundtruth attendance estimates, however – for the same reason – it is also the best scenario for the actual use of the proposed technique.

Figure 12 illustrate results for a set of “unstructured” events. We fit the linear regression by using “structured” events (football matches in stadiums) happening in the same city and we searched the Web for reported attendance estimates and use them as ground truth. In this case results are worse than in the previous case, obtaining 22% median %error. On the one hand, this is due to the fact that these events tends to be smaller than “structured” ones, thus making the attendance estimation task inherently more difficult. On the other hand, the linear regression is trained for larger “structured” events, thus it can be a less effective fit for these events. Finally, as groundtruth estimate for this events is weaker, the fact that a number of events have an estimated attendance *lower* than the groundtruth might be




Figure 10: Attendance estimation results. (top row) correlation plot with different kinds of regressions. (left) linear, (center) piecewise, (right) range. The shaded area represents confidence interval for the regression outcome. (bottom row) r^2 , mean/media absolute and percentage errors for the different regressions.

also interpreted as the fact that the estimates reported in the news (on the Web) are inflated.

5. KNOWLEDGE OF MULTIPLE EVENTS

In all the previous algorithms and experiments we considered events in isolation: we tried to estimate the attendance to an event *without* any information about other events happening in the same place. On the contrary, if we know a number of events that happened in a given place, we can adopt a different procedure to estimate the radius of the event area.

In particular, we try to estimate the area associated to a given placemark (e.g., a stadium), and all the events happening in there will be associated to the same event area. This procedure updates the procedure described in Section 3.2 STEP 2. The idea is that, instead of weighing each possible radius by how extreme values (z-score) it produces, we weight each radius by how many events it is able to identify as outliers. This is basically the procedure in Figure 6 in which we count the number of events. More in detail the process to identify the event radius is the following:

STEP 1. The same as in Section 3.2

STEP 2.

1. For each event area, we identify a number of events happened in there (this will serve as a sort of “training” set).
2. We consider the z -score computed in STEP 1 for a time frame encompassing all the events in the training set.
3. We identify outliers in the z_i values as those points with a value greater than 3. We then count the number of outliers that happens to be at the same day and time of events in the “training” set. The result is that for each value r_k we have the number of events being identified e_k .
4. The final value of the radius $best_r$ is the average of the r_k values weighted by the number of identified events e_k

See a more formal description in algorithm 4. On the one hand this approach tends to be more robust in that radius parameters are choose to detect the largest number of events. On the other hand, it is less flexible in

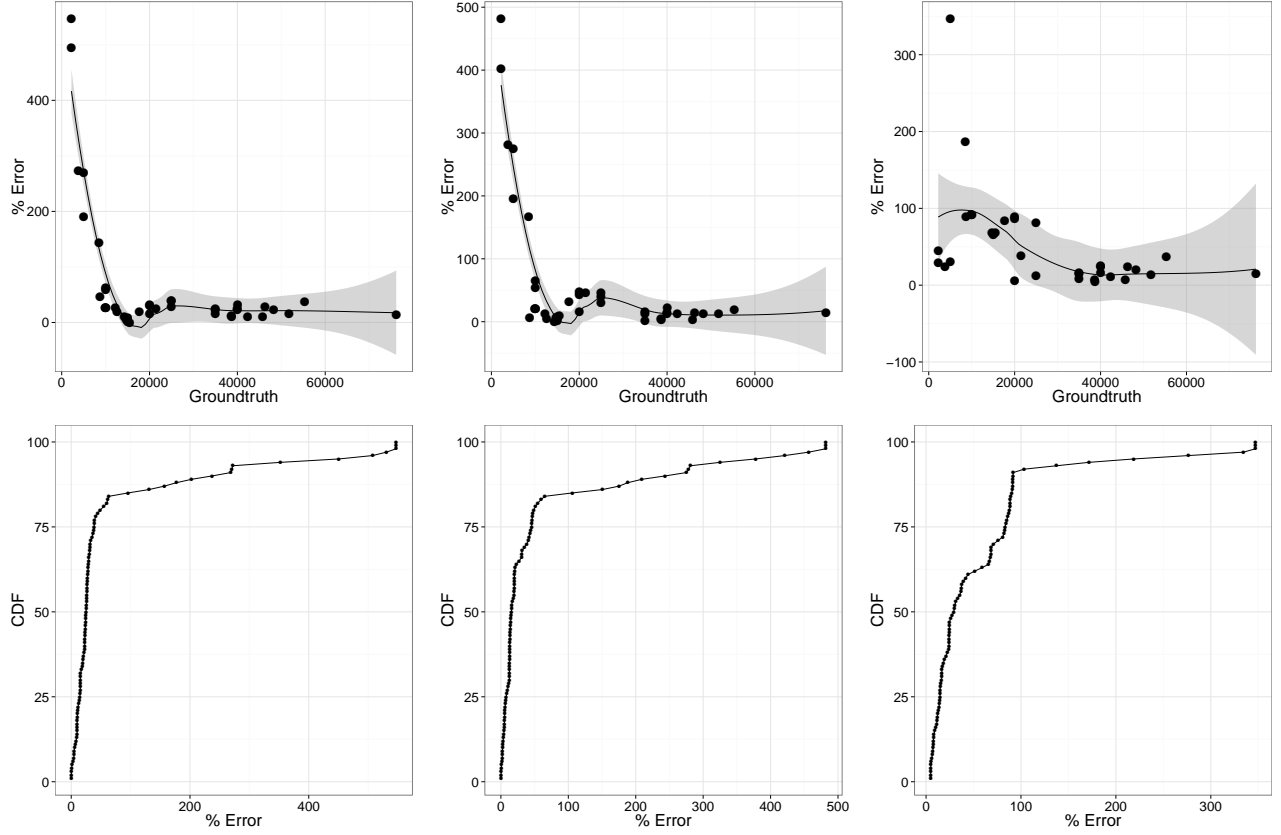



Figure 11: Attendance estimation errors. (top row) error distribution for different regressions: (left) linear, (center) piecewise, (right) range. (bottom row) error cumulative density function. It is possible to see that error distribution is highly skewed. Our approach is effective for large events, while it has considerable errors for small ones.



Piecewise R.	
Mean abs. error	7307
Median abs. error	4422
Mean % error	33%
Median % error	22%

Figure 12: Results obtained by using piecewise linear regression for “unstructured” events

```

Data:  $z[]$ 
Result:  $bestR$ 
forall the  $r_k \in [r_{min}, r_{max}]$  do
     $x_k = countUsers(cdr[], ec, r_k, st, et)$ 
     $e_k = \text{number of event days with } z_k > 3$ 
     $bestR = \frac{\sum_k r_k \cdot e_k}{\sum_k e_k}$ 
end

```

Algorithm 4: Multiple events. Radius Extraction - Step 2

that it associates a single radius to a given place without the flexibility of enlarging the radius for larger events in the same place. Figure 13 illustrates the results of the estimation approach with the radii computed in this way and adopting piecewise linear regression. (left) correlation plot, (center) % error distribution, (right) % error CDF. Overall, the knowledge of multiple events further improves the results: $r = 0.95$, $r^2 = 0.9$, median absolute error = 4160, median % error = 10%.




Figure 13: Results obtained by adopting the radii computed with the knowledge of multiple events. (left) correlation plot, (center) % error distribution, (right) % error CDF

6. RELATED WORK

The application potential of estimating the number of people present in specific parts of the city at specific times led to the development of a number of approaches and researches to tackle the issue. Such estimates can in fact provide useful information to the local government and to the event organizers to plan, manage and respond to the event. Also the advertisement industry would get notable information from such data, in that it is possible to measure how many people were able to see a given advertisement, understand where they come from, their habit, etc.

People counts, surveys, and other traditional methods to identify the presence of visitors and tourists in a city are often expensive and result in limited empirical data. Similarly, the exploitation of land-use (e.g. density of hotels) and census data provides only a static perspective on city dynamics. The lack of data presents particular difficulties given that most cities – though they may aim at providing advanced services – have limited human, technical, and financial resources. Today, thanks to the emergence of ubiquitous technologies, new data sources are available

Fueled by the “recent” availability of telecoms’ CDR data, a number of researchers try to automatically identifying events happening in the city and estimating the number of people attending the event.

The works in [16, 17] present an approach to estimate the attractiveness of events happening in the city from the combination of cellular network activity and other information sources. They try to estimate the location of cellular network traffic and to use it as a proxy of the number of people in that area. However, these methods can identify daily trends and outliers, but they can not estimate the actual number of people.

The work presented in [8, 18] presents another ap-

proach to analyze people attendance to special events on the basis of CDRs coming from the AirSage (www.airsage.com) platform. In this work, they segment users’ traces to identify those places where a user stops. If this place, coincides with the place of the event and if the duration of the stop is at least the 70% of the duration of the event, the user is classified as attending the event. On this basis they are able to analyze the attendance to specific events. However, they claim: “*Estimating the actual number of attendees is still an open problem, considering also that ground truth data to validate models is sometime absent or very noisy*” and do not perform quantitative analysis in this direction.

The work in [14] is very interesting and closer to our approach. They use a Bayesian model to localize the source of CDRs. Then, they compute the probability p of each user to participate an event as $p = p_1 \cdot (1 - p_2)$. Where p_1 is the fraction of time in which the user is in the event area at the event time. p_2 is the fraction of time in which the use is in the event area at other times. Finally, they use an outlier detection mechanism (based on a z-score) to classify users as participants to an event. Unfortunately, they use the approach only to identify an event and not to estimate the actual attendance.

A similar approach to identify events is reported in [3]. In this work, authors apply an outlier detection mechanism to aggregated cell network data (i.e., erlang measurements). Events are associated to overcrowded or suddenly underpopulated areas.

In conclusion of this section, while some works propose approaches to detect and analyze events happening in the city by using the data from cellular network, the problem of actually estimating the number of attendees is largely unexplored. In particular, to the best of our knowledge, there are not published results of the accuracy of attendance estimation using CDRs.

7. CONCLUDING REMARKS

In this work we propose an innovative methodology to estimate the number of attendees to events happening in the city from cellular network data. We evaluate our approach in 43 events ranging from football matches in stadiums to concerts and festivals in open squares. Comparing our results with the best groundtruth data available, our estimates provide a median error of less than 15% of the actual number of attendees.

While the obtained results are very encouraging, there are a number of research directions that could improve the presented work:

- Of course, running experiments on other, more diverse, events would better validate our results.
- Our work has been mainly driven by experiments. A better theoretical framework for our modeling (especially with regard to the event area estimation) could provide further ideas for improvement.
- A deeper analysis of the trajectories of individual users could provide a more fine grained localization of CDRs, thus leading to a better estimate of the user's presence in the event area [6]

Despite the above limitations, to the best of our knowledge, this is the first work providing a practical and accurate way of estimating the number of attendees to events happening in the city from cellular network data.

8. REFERENCES

- [1] F. Calabrese, C. Ratti, M. Colonna, P. Lovisolo, D. Parata, Real-time urban monitoring using cell phones: a case study in rome, *IEEE Transactions on Intelligent Transportation Systems* 12 (1) (2011) 141–151.
- [2] L. Ferrari, M. Mamei, Discovering city dynamics through sports tracking applications, *IEEE Computer* 44 (12) (2011) 61–66.
- [3] L. Ferrari, M. Mamei, M. Colonna, Discovering events in the city via mobile network analysis, *Journal of Ambient Intelligence and Humanized Computing* 5 (3) (2014) 265–277.
- [4] R. Becker, R. Caceres, K. Hanson, S. Isaacman, J. M. Loh, M. Martonosi, J. Rowland, S. Urbanek, A. Varshavsky, C. Volinsky, Human mobility characterization from cellular network data, *Communications of the ACM* 56 (1) (2013) 74–82.
- [5] F. Zambonelli, Toward sociotechnical urban superorganisms, *IEEE Computer* 45 (8) (2012) 76–78.
- [6] I. Leontiadis, A. Lima, H. Kwak, R. Stanojevic, D. Wetherall, K. Papagiannaki, From cells to streets: Estimating mobile paths with cellular-side data, in: *International Conference on emerging Networking EXperiments and Technologies (CoNEXT)*, Sydney, Australia, 2014.
- [7] N. Lathia, V. Pejovic, K. Rachuri, C. Mascolo, M. Musolesi, P. Rentfrow, Smartphones for large-scale behaviour change intervention, *IEEE Pervasive Computing* 12 (2013) 66–73.
- [8] D. Quercia, G. D. Lorenzo, F. Calabrese, C. Ratti, Mobile phones and outdoor advertising: Measurable advertising, *IEEE Pervasive Computing* 10 (2) (2011) 28–36.
- [9] D. J. Mir, S. Isaacman, R. Caceres, M. Martonosi, R. N. Wright, Dp-where: Differentially private modeling of human mobility, in: *IEEE International Conference on Big Data*, Santa Clara (CA), USA, 2013.
- [10] A. Basu, A. Monreale, J. C. Corena, F. Giannotti, D. Pedreschi, S. Kiyomoto, Y. Miyake, T. Yanagihara, R. Trasarti, A privacy risk model for trajectory data, in: *International IFIP Conference on Trust Management*, Singapore, 2014.
- [11] M. Rahnema, Overview of the gsm system and protocol architecture, *IEEE Communications* 31 (4) (1993) 92–100.
- [12] R. Caceres, J. Rowland, C. Small, S. Urbanek, Exploring the use of urban greenspace through cellular network activity, in: *Workshop on Pervasive Urban Applications (PURBA)*, Newcastle, UK, 2012.
- [13] M. Ulm, P. Widhalm, Properties of the positioning error of cell phone trajectories, in: *NetMob*, Boston (MA), USA, 2013.
- [14] V. A. Traag, A. Browet, F. Calabrese, F. Morlot, Social event detection in massive mobile phone data using probabilistic location inference, in: *International Conference on Social Computing*, Boston (MA), USA, 2011.
- [15] M. Gonzalez, C. Hidalgo, A. Barabasi, Understanding individual human mobility patterns, *Nature* 453 (2008) 779–782.
- [16] F. Girardin, A. Vaccari, A. Gerber, A. Biderman, C. Ratti, Quantifying urban attractiveness from the distribution and density of digital footprints, *International Journal of Spatial Data Infrastructure Research* 4 (2009) 175–200.
- [17] J. Neumann, M. Zao, A. Karatzoglou, N. Oliver, Event detection in communication and transportation data, *Pattern Recognition and Image Analysis* 7887 (2013) 827–838.
- [18] F. Calabrese, F. Pereira, G. Lorenzo, L. Liang, C. Ratti, The geography of taste: Analyzing cell-phone mobility and social events, in: *International Conference on Pervasive Computing*, Helsinki, Finland, 2010.